

THE CONVENTIONAL WISDOM OF BEHAVIOR ANALYSIS

J. E. R. STADDON

DUKE UNIVERSITY

To those outside the field of behavior analysis, its philosophical assumptions and the strictures they place on experiment and theory often seem narrow and doctrinaire. To behaviorists, experimental work outside often seems equally constrained, compressing the rich and sometimes idiosyncratic behavior of individual subjects into impoverished statistical summaries that filter out anything of real interest. Theory by nonbehaviorists—cognitive theory, for example—usually appears to behaviorists to be vague or excessively elaborate, often full of mentalistic concepts and only weakly linked to behavior and the subject's actual experience. I can sympathize with both positions. But as someone closer to the behaviorist camp, I am more alarmed by legitimate criticisms of the experimental analysis of behavior than by the weaknesses of its competitors.

Marc Branch's editorial in the January 1992 issue of the *Journal of the Experimental Analysis of Behavior (JEAB)* provides an excellent focus for some of these concerns because it is a brief and elegant presentation of the *conventional wisdom* (J. K. Galbraith's felicitous phrase) of the experimental analysis of behavior. Branch summarizes the standard practice of radical behaviorists, making points most would accept without question. Yet I will argue that many of these points, either on their face or as they are usually understood, are far from self-evident. One or two are highly debatable. I here take up four of these points.

The following passage highlights the first two:

"Environment-based" theorizing, developed in psychology expansively by Skinner . . . but having its roots much earlier (cf. Darwin, 1872/

1962; Mach, 1883/1960 . . .), may be contrasted with "organism-based" interpretations. . . . A good analogy is to Newtonian mechanics. . . . I hope *JEAB* will continue to emphasize environment-based explanation, not because it is necessarily "right" but because it is an approach with advantages worth exploring. . . . (Branch, 1992, p. 1)

Let's look at two implications of this passage, and the paragraph from which it is taken: (a) that "environment-based" theorizing is distinctive of Darwin and sanctioned by Mach, and (b) that Newtonian mechanics is an appropriate model for theorizing about behavior.

Precedents for Environment-Based Theorizing

Darwin dealt primarily with environmental determinants only because behavior was not his primary object of study. Even in his work on phylogeny, however, he was careful to distinguish between *present* and *past* environments—a distinction often blurred in the usual term *history*: "For natural selection acts by either now adapting the varying parts of each being to its organic and inorganic conditions of life; or by having adapted them during past periods of time . . . (Darwin, 1872/1951, p. 217, my emphasis). In the experimental analysis of behavior, *history* almost invariably refers only to present, or recent, environments—not to events in the remote past.

When Darwin did deal with behavior, his theorizing was more "organism based" than might be supposed. For example, he discusses the behavior of ants carrying their cocoons away from the nest in the following way:

The ants carrying the cocoons did not appear to be emigrating. . . . But when I looked closely I found that all the cocoons were empty cases. . . . Now here I think we have one instinct in contest with another and mistaken one. The first instinct being to carry the empty cocoons out of the nest. . . . And then came in the contest with the other very powerful instinct of preserving and carrying their cocoons as long as possible; and this they could not help doing although the cocoons were empty. According as the one or other instinct was the stronger in each individ-

I thank members of the Learning and Adaptive Behavior Group for vigorous discussions of these issues and helpful comments on earlier versions of this note. Research was supported by grants to Duke University from NSF and NIMH. Address correspondence to J. E. R. Staddon, Department of Psychology: Experimental, Duke University, Durham, North Carolina 27708-0086.

ual ant, so did it carry the empty cocoon to a greater or less distance. (Darwin, 1891, p. 370)

This informal account could be made precise by defining the two competing action patterns: carrying the cocoons out of the nest and holding on to the cocoons as long as possible. Given hypotheses about the effects of various manipulations on the relative strengths of these two "instincts," as Darwin calls them, predictions about the distance the cocoons should be carried could be made and tested. A theory of this sort surely qualifies as organism based. Yet even though it gives us no clue as to the provenance of the "instincts," it seems to me a perfectly respectable scientific procedure. The point is that even Darwin was not averse to postulating entities that are not directly observable, if a puzzling phenomenon ("Why do the ants take the cocoons so far from the nest?") is thereby made more comprehensible.

Branch's brief reference to physicist Ernst Mach is an interesting example of Mach's stature in experimental psychology. Mach's influence has been largely indirect, through the translations and interpretations of S. S. Stevens (e.g., 1951), E. G. Boring (e.g., 1950), and their students in experimental psychology, and in the experimental analysis of behavior through Mach's influence on Skinner during his graduate days at Harvard (cf. Marr, 1985; Smith, 1986). Mach's prestige in psychology is high, and his views (as seen through the interpretations of Boring, Skinner, and others) are rarely questioned.

In physics Mach's brilliance is universally acknowledged. But his philosophy of science (which is of course the main import from his work to psychology) is not so widely admired. For example, Einstein commented:

Mach's weakness, as I see it, lies in the fact that he believed more or less strongly that science consists merely of putting experimental results in order; that is, he did not recognize the free constructive element. . . . He thought that somehow theories arise by means of *discovery* and not by means of *invention*. ("Einstein," 1991, p. 35)

Einstein's point is that vulgar Baconianism is wrong: Not all theories can be arrived at simply from orderly arrangements of data.¹ For

every Mendeleev and his period table (a good inductive theory) there is a Faraday, a Mendel, or a Kékulé who arrives at field theory, genes, or the benzene ring by a leap of the imagination.² As any science matures and orderly data (inductively gathered) begin to accumulate, the creative element—the ability to invent correct theories that are not obvious from any likely arrangement of data—becomes essential to further progress. Unfortunately, the uncritical importation into experimental psychology, and particularly into behavior analysis, of Mach's proinductive prejudice (and Skinner's antitheoretical one) has stunted the development of theory in our area and fostered unnecessary antagonism between behaviorism and approaches that are more catholic (perhaps too catholic!) in their theoretical explorations.

Newtonian Mechanics

Is Newtonian mechanics an appropriate model for behavior theory? Its attractions are obvious. It deals only with observables like mass, length, time, and higher derivatives of these. We like to think that measures like response and reinforcement rates share this kind of rigor. But the problem with Newtonian mechanics as a model for behavior theory is that it is entirely *ahistorical*. Given the initial conditions of the system (and in Newtonian physics these are all directly observable quantities), the system's future behavior (reactions to new forces, etc.) can be predicted. Given the same initial conditions, the future behavior must be the same.

The problem for behavior is that the same set of observables—response rates, preferences, or whatever—at one time (t_1) denote a different *system state* than an identical set at another time (t_2). How do we know that the system state is different? Because the same set of experimental manipulations (extinction, for example) may produce different results (e.g., greater or less resistance to extinction) if ap-

course Mach and Skinner felt little need for theory at all. (Skinner's book, 1969, *Contingencies of Reinforcement*, is subtitled *A Theoretical Analysis*, but the "theory" is far from the kind of formal system favored by other learning theorists.)

² I use the term *imagination* here not to connote some mysterious or inexplicable process, but simply to note that these discoveries did not follow in any obvious or inevitable way from the data they were used to explain.

¹ Bacon in fact got it right. He never actually argued that science is nothing but orderly fact gathering. And of

plied at t_1 than at t_2 . The point is that the future behavior of a historical system cannot be predicted from observables alone. Simple Newtonian mechanics is not an appropriate model for historical systems and therefore should not provide a model for behavior theory.

One objection to this conclusion is to point to examples of static, empirical laws, like Weber's law or the matching law, that are not historical. But, the animal that matches today is nevertheless not the same animal that matched last week, and any science that aspires to understand how matching comes about must take historical dependence into account. Unless we are content to remain forever at the level of static principles, and thus abandon any hope of understanding the process of learning, something beyond the Newtonian model must be found. More on choice in a moment.

A second objection is more technical. It consists in redefining *initial conditions* in such a way as to take in a substantial chunk of the system's past history. It is a valid objection, in the sense that given a known deterministic system, and given a sufficient set of historical observations, it is possible to define the current state of the system in such a way as to fulfill the Newtonian property that future behavior can be predicted perfectly. But this objection simply makes my final point that the state of a system and its history are transforms of one another. More on this point also in the final section.

There is one other point I want to discuss before getting to the assumed contradiction between "environment-based" and "organism-based" accounts of behavior: the relation between language and theory in technical descriptions of behavior.

Language

Branch (1992) writes, "A calculus based in mathematics has been useful in science partly because of its lack of ambiguity. A calculus based in verbal behavior also should be as unambiguous as possible" (p. 2). Mathematical manipulations are indeed unambiguous, but this is not the main reason mathematics is so useful in science. Mathematics has been useful not so much because it is precise as because it is the language of *theory*. Even the most elementary mathematical theory (the Newtonian derivation of the period of a pendulum from

its length, for example) would be essentially impossible to express in words (cf. Staddon, 1984). Mathematics is essential not so much because it is precise as because it permits the derivation of predictions that would be impossible to arrive at through verbal arguments alone.

Moreover, mathematical theory is precise only in its manipulations of symbols. The meaning of those symbols, the measurements necessary to assign values to them, rest on the verbal descriptions. Fortunately, the theory itself constrains and guides the definitions of the terms that enter onto it. The terms in the equations are operationally defined, but the definition is often implied by the theory. This constraint is important because we can define things in any way we please. Without an implicit or explicit theory, there is no reason to prefer one operationally defined term to another. (Not all technical terms are theoretical, of course. The definitions of structures and procedures are simply descriptive and arouse no controversy. The problem is terms with theoretical overtones.)

Consider *force*, for example. Newtonian law says that $F = ma$, which suggests that we define unit force as that which will accelerate one unit of mass one unit of acceleration. But if the law of nature were $F = mva$, where v = volume, this definition of force would be useless, because it would correspond to different physical quantities for different volumes. The proper definition would have to include volume. The point is that the law specifies the verbal definition, *not* the reverse. Hence, if we do not know the law, *we have no business being rigid about verbal definitions!* Insisting on a term before we have an accepted theory is putting the cart before the horse.³

The problem of defining concepts in advance of theory is well known in the natural sciences. Peter Medawar quotes Herbert Spencer, who wrote as follows on the limited virtues of precision when fundamental knowledge is lacking:

³ Sometimes the appropriate definition is obvious from "orderly functional relations" (e.g., between weight and imparted acceleration). But the history of physical science shows that the development of a formal theory is usually necessary before definitions attain universal acceptance. Acceptable definitions rarely rise from empirical regularities alone.

A preliminary conception, indefinite but comprehensive, is needful as an introduction to a definite conception. A complex idea is not communicable directly, by giving one after another its component parts in their finished forms; since if no outline pre-exists in the mind of the recipient these component parts will not be rightly combined. Much labour has to be gone through which would have been saved had the general notion, however cloudy, been conveyed before the distinct and detailed delineation was commenced. (Medawar, 1967, p. 43)

Many other philosophers of science have made similar points. Mach enjoined against formulations involving "greater precision than fits the needs of the moment." Even Skinner, in later life so concerned with correct behavioral language, in early days warned against making "a fetish of exactitude" (both these examples from Smith, 1986, p. 270). See also philosopher of biology David Hull's (1988) illuminating discussion of "weasel words" in science. But the obsession with premature precision, a legacy of the mistaken positivism of P. W. Bridgman, still bedevils parts of psychology—the experimental analysis of behavior especially.⁴

This is not to excuse vagueness, of course. People should be precise about the words they use. But there is absolutely no reason to be compulsive about the particular words in areas in which understanding is still imperfect, so long as we are clear what is intended. For example, *JEAB* used to insist on the word *reinforcement* rather than *reward*, or even *food*. But even now we cannot be certain that all the phenomena that fit the conventional definition of reinforcement (response-rate increase caused by contingent reinforcer presentation) in fact belong together, because we still have no consensus on the underlying process (or processes). Hence, to insist on the word *reinforcement* as opposed to *reward* or simply *food*

delivery is to imply greater precision than the state of knowledge warrants. On the other hand, most *JEAB* readers are alert to the improper connotations of the vernacular word *reward*, so that few will read into it anything more than the delivery of food to a hungry animal.

The kind of "behavioristic correctness" that is represented by insistence on particular terms (rather than simply on clarity) can easily become scientifically counterproductive. It has the beneficial effect of forcing novices to think about the words they use. But it has the inverse defect of allowing experts to write without thinking. No doubt the possession of a shared (and preferably esoteric) vocabulary enhances group cohesion—the concluding sentence, "May your reading be consequential," of Catania's (1991) recent editorial conveys something of this cozy feeling. But a technical vocabulary that is not firmly grounded in theory or descriptive utility also limits speculation, alienates outsiders, and marginalizes the field.

Environment-Based Versus Organism-Based Explanations

This is the kernel of Branch's editorial, and the main point on which behaviorists and cognitivists differ. Yet the opposition between organism-based and environment-based theories is only a difference of emphasis, as I will try to show. The easiest way to do this is by example.

The example is taken from the work of Derick Davis in my laboratory on choice between random-ratio schedules (the so-called "two-armed bandit" problem). The essentials of the situation are that hungry pigeons choose between two response keys that deliver food reinforcement according to the probabilistic schedules. Each day, the probability is one in eight that a peck to one of the two keys will produce food. The "hot" key varies from day to day, as specified by the experimenter.

The upper plot in Figure 1 shows the kind of data we get (Davis, 1991). The solid line shows percentage "correct" (i.e., responses to that day's S+) for a single pigeon each day. The dashed line shows learning rate, parameter *a*, computed as described in the legend. Low values of parameter *a* correspond to rapid learning. These data show three effects of reversal learning: First, there was improvement in reversal performance across successive daily

⁴ This is not the place for lengthy examples, but the preoccupation with definition in the absence of theory begins in the experimental analysis of behavior with Skinner's (1935) paper on the operation analysis of psychological terms. More recent examples include Logan and Ferraro's (1978) positivistic text, various attempts to codify psychological terms (e.g., Verplanck, 1957, and unpublished) or experiments (Logan, unpublished), and the several sections on "vocabulary" in Catania's (1984) popular text. See also Catania (1989) for a discussion of the role of language in the experimental analysis of behavior.




Fig. 1. Top: the entire course of a discrimination-reversal experiment with a single pigeon. The solid line shows the proportion of responses each day on the key that currently produced food, when those conditions were alternated every day, every 2 days, or every 4 days, as shown by the staggered letters at the top: L = only left responses reinforced; R = only right responses reinforced (from Davis, 1991). The dashed line shows the rate of learning from Day N to Day $N + 1$, computed from the formula $a = [S(N + 1) - X(N + 1)] / [S(N) - X(N + 1)]$, where S is the proportion of right responses [$=R / (R + L)$], and the reinforcement asymptote $X = 0$, for left reinforcement and 1 for right. This formula is derived from the linear integrator model for reversal learning (Davis & Staddon, 1990). Bottom: simulation of these data by the cumulative effects model with initial conditions 1,000, 2,000. See Davis et al. (1993) for more details.

reversals. This is shown both by daily improvement in percentage of correct responses (solid line) and by the increased learning rate (reduced value of parameter a) across the series of daily reversals (single alternation) in the left half of the top graph. Second, there was a change in learning rate (parameter a) as a function of frequency of reversal: a was lowest in daily reversal and highest when reversal was every 4 days. Third, there was a change in learning rate within and between blocks in the 2- and 4-day reversal conditions: The computed a value increased within each block and then decreased between blocks.

What would constitute an explanation for these results? At one time, improvement in successive reversals was attributed to the development of a "reversal learning set." We now know that this is not much of an explanation. Branch (personal communication) has suggested, tongue-in-cheek, that an "organism-based theorist might propose a 'behavioral alternation device' (BAD!)" that has the required properties. Properly chosen, such a BAD might look like an explanation of the data I have described and might even suggest predictions about future experiments. But the provenance of the BAD, its origin in pigeon genes or in the training procedure, would necessarily not be part of the theory. I would be reluctant to rule out such an account completely—any model that can predict powerfully deserves to be taken seriously (cf. Staddon & Bueno, 1991)—but in the abstract this theoretical strategy looks unpromising.

What is the alternative? At least one behavior analyst has suggested that on daily reversals the pigeon learns to use the initial reinforcement each day as a discriminative stimulus for choice the rest of the day. There is some support for this idea (Williams, 1976), but its predictive power is limited and it provides no obvious explanation for the differences between single, double, and quadruple alternation. This account also shares at least one unfortunate property with the BAD: It provides no principled basis for identifying what will be the important discriminative stimulus. Why choose the preceding reinforced response? After all, the animal could do even better on daily alternation if it used the identity of S- on Day N as a cue for S+ on Day $N + 1$, but pigeons never do so. The theory is

silent on why the animal selects the less effective S^D over a more effective one. Like BAD, this is a theory of "what is learned," not of how it is learned.

Cumulative effects (CE) model. Here is an alternative that I offer, despite its imperfections, because it illustrates the arbitrariness of the dichotomy between environment- and organism-based theories. My colleagues and I (Davis, Staddon, Machado, & Palmer, 1993) have proposed the following process as an explanation for these reversal-learning data. We suggest that the pigeons choose response by response which key to peck on the basis of a quantity we call V : The key with highest V is the one pecked. V is nothing but reinforcement probability, computed in a slightly novel way. For each key, $V = (R + R_0)/(N + N_0)$, where R is the total number of reinforcements and N the total number of responses *since the beginning of the experiment*. R_0 and N_0 are simply initial conditions: the values for R and N at the beginning of the experiment. There is a V value for each choice, and the choice with the highest current V value is always the one chosen; that is, the process, which we call the *cumulative effects (CE) model*, is a version of momentary maximizing (Shimp, 1966).

If one simulates this process, it turns out to provide a pretty good facsimile for the data in the top panel of Figure 1. That is, it shows improvement in percentage correct across successive reversals, plus the other effects of double and quadruple alternation that I have described. It also explains a number of other properties of behavior in situations like this (matching, for example). A typical simulation is shown in the bottom plot in Figure 1.

Now the question: Is this an organism-based or an environment-based explanation? My answer is "both, depending on how you describe it." It is organism based if you focus on the fact that the behavior of the (model) organism depends on an internal state defined by four variables (in a two-choice situation): the values of $R + R_0$ and $N + N_0$ for each choice (note that these values cannot be estimated from single-session average data). But it is equally an environment-based explanation if you focus on the fact that these state variables can be computed from the animal's past history, although the fact that the initial conditions must be derived indirectly from the data means that the

state even of this very simple model is not isomorphic with directly measurable quantities.

Moreover, the model "state" here is both more and less than the organism's (model's) past history. It is more than past history because any given organism has only one history, but the model summarizes the effects of the infinite set of past histories that all end up with the same four-place vector of R and N values. In other words, the model tells us what other experiences the organism *might* have had that would have left it the same in terms of all future behavior. This is the idea of *equivalent history*, to which I will return in a moment.

The model state is also more than a given history in the sense that it includes initial conditions, which are not themselves observable or directly measurable, but must be estimated from data. But the model is less than history in the sense that once we know the model, we can test it at any time so as to estimate the values of the four state variables. Given the model, plus the results of a test, we can therefore estimate the state—hence the model's future behavior—without knowing the details of its past history. This is a direct, practical benefit for an accurate model over bare knowledge of historical data.

Despite the arguments for the value of the "state" idea, the CE model would probably be regarded as environment based by most people, because its state is so closely linked to observable behavior. But this model is only a special case of a much more complex class of models of the same general type. Let me briefly describe a slightly more general version that most people would probably term organism based.

One perfectly reasonable objection to the CE model is that because the V values simply reflect totals, the influence of a reinforcement just received is no greater than the influence of a remote reinforcement, received yesterday, or last week. Even though the model's predictions are pretty good, this assumption seems implausible. The assumption can be modified in various ways. One is to use some kind of temporal weighting, such as Killeen's (1981) exponentially weighted moving average. This keeps us still pretty close to the raw data, but in our simulations it doesn't work very well.

Cumulative trace (CT) model. Another pos-

sibility is as follows. Recall that the original model defines a quantity V that is simply a ratio of cumulated responses and reinforcers for each choice. Thus, the numerator of this ratio is $r_1 + r_2 + \dots + r_N$, where r denotes a reinforcer magnitude and the subscripts are just the first, second, \dots N th reinforcers. Now, suppose we accept that older reinforcers have diminishing effects. This suggests that the contribution of each reinforcer should decrease with the time elapsed since it was delivered. Thus, the numerator should be something like $r_1 f(t - t_1) + r_2 f(t - t_2) + \dots + r_N f(t - t_N)$, where t_i is the absolute time of occurrence of the i th reinforcer, t is the current time, f denotes a decreasing function of age (i.e., $t - t_i$) such as $e^{-a(t - t_i)}$, and r_i is just the reinforcer magnitude. A similar function is used for the responses in the denominator. With this change, the longer ago a response or a reinforcement occurred, the smaller its contribution to the total.

I don't know whether this model is in fact better than the CE model. We do know that it can mimic time-based spontaneous recovery, which is not explainable by the response-level CE model. But there are other, computationally simpler, models that seem worth exploring first. The virtue of the CT model in the present context is that its family resemblance to the CE model permits the following argument.

This model is simply a generalization of our original CE model. In the original, we just picked a particularly simple form for f , namely $f(t - t_i) = 1$, that is, a constant value for f independent of $t - t_i$. Nevertheless, despite the added computational complexity of the second model, the two models are of exactly the same form. Yet the second is in fact a *memory trace* model, because $f(t - t_i)$ represents a declining influence of the i th reinforcer with time. Moreover, the CT model adds at least one parameter, representing the rate of decay of the "trace" (two parameters, if we assume that the reinforcement and response traces decay at different rates). Because of these new parameters, it is much more difficult to estimate the state of the CT model from behavioral data. The CT model shares some properties with Hullian trace theory, and most people would therefore consider it an organism-based model. Yet, I would argue, despite the much greater computational complexity of the CT model, that

it is no less environment based, and no more organism based, than the CE model.

The point is that the environment-based versus organism-based distinction is often impossible to make in practice.

States and equivalent histories. The same point can be made more abstractly. There cannot be any real opposition between environment- and organism-based explanations, for a very fundamental reason: The "internal state" of any black-box system cannot be known *except* through knowing its history. This is a matter of logic, not of data or definition. Talking about internal states is thus just a shorthand way of talking about sets of equivalent (behavioral) histories—equivalent in the sense that the future behavior of the system is the same following any of the histories in the set (cf. Minsky, 1967; see also Himeline, 1990; Staddon, 1973).⁵ Thus, for the CE model, all histories that end up with the same totals for $R + R_0$ and $N + N_0$ are equivalent, in the sense that the future behavior of the model will be indistinguishable.

Notice that because the state of the CE model is not identifiable from session-average data, the model is not subject to the Newtonian constraint I discussed earlier: It shows effects of past history that may not be visible in currently measured behavior. Similar session-average data early and late in training do not, therefore, imply similar behavior when conditions are changed.

The logical equivalence between a state and a set of equivalent histories means that there cannot be any real conflict between organism- and environment-based explanations for behavior. Any environment-based theory can be rephrased as an organism-based theory, although the converse is not true. A valid criticism of some cognitive theories is that states are postulated without any clear specification of the historical data necessary to identify them. This asymmetry may lead some to argue for the intrinsic superiority of the environment-based approach, but there are counterarguments I will not go into here. For the CE model, the definition of the model in environment-based terms is just about as simple as its

definition in terms of internal state. But for the CT model, the environment-based definition is probably easier to understand, even though most people would probably categorize the model as organism based.

If these two types of models are intertranslatable, as I have argued, what then is the basis for the vigorous disagreement between behaviorists and cognitivists? This question takes us into the realm of conjecture and goes well beyond what can be treated in a short paper. Evidently most cognitivists prefer theory to data, and give some priority to mentalism and intuition (compared to induction from data) in the invention of theories. Behaviorists clearly expect much more in the way of empirical support for any theoretical speculations, and give little weight to the intuitive plausibility of a hypothesis. Moreover, many behaviorists simply do not accept that the ultimate objective of psychological science is theoretical: "behavior in its own right" is still a potent talisman. Nevertheless, the logic is as clear for behaviorists as for cognitivists. If we disavow an interest in either physiology or mind reading, then we cannot know anything about internal states except through the study of particular histories. And if the finite effects of an infinite set of possible histories are to be reduced to some kind of order, groupings of histories that are in effect states are unavoidable. Behaviorists would therefore do well to acknowledge the utility of internal, albeit historically defined, states and judge a theory on its explanatory merits rather than on whether or not it is "environment based."

Conclusion

It is time to move beyond an essentially atheoretical and ahistorical behaviorism to a theoretical and historical behaviorism that recognizes the obvious: Histories do not act in a vacuum; the organism is not simply the passive confluence of forces, like a Ouija® board pushed by intoxicated seancers. Skinner (1972), despite his protestations, was not just the empty locus at which a poem "happened" but a unique living being changed by his experiences and operating upon them in creative ways. Histories *change the organism*, and experimentally based state models are the only way we have to understand the nature of the changes and the nature of the mechanism that allows them to occur.

⁵ A Skinnerian reader might want to consider this definition of *state* as a class of histories as merely an extension of Skinner's (1935) class-based definitions of stimulus and response (see also Staddon, 1967).

REFERENCES

- Boring, E. G. (1950). *A history of experimental psychology*. New York: Appleton-Century-Crofts.
- Branch, M. N. (1992). On being narrowly broad. *Journal of the Experimental Analysis of Behavior*, **57**, 1-4.
- Catania, A. C. (1984). *Learning* (2nd ed.). Englewood Cliffs, NJ: Prentice-Hall.
- Catania, A. C. (1989). Speaking of behavior. *Journal of the Experimental Analysis of Behavior*, **52**, 193-196.
- Catania, A. C. (1991). Contingent reviews. *Journal of the Experimental Analysis of Behavior*, **56**, 591-598.
- Darwin, C. (Ed.). (1891). *Life and letters of Charles Darwin*. New York: D. Appleton.
- Darwin, C. (1951). *The origin of species*. London: Oxford University Press. (reprint of the 6th edition, 1872)
- Davis, D. G. S. (1991). *Probabilistic choice: Empirical studies and mathematical models*. Unpublished doctoral dissertation, Duke University.
- Davis, D. G. S., & Staddon, J. E. R. (1990). Memory for reward in probabilistic choice: Markovian and non-Markovian properties. *Behaviour*, **114**, 37-64.
- Davis, D. G. S., Staddon, J. E. R., Machado, A., & Palmer, R. G. (1993). The process of recurrent choice. *Psychological Review*, **100**, 320-341.
- Einstein the inventor. (1991, Winter). *Invention and Technology*, pp. 34-39.
- Hineline, P. N. (1990). The origins of environment-based psychological theory. *Journal of the Experimental Analysis of Behavior*, **53**, 305-320.
- Hull, D. L. (1988). *Science as a process*. Chicago: University of Chicago Press.
- Killeen, P. R. (1981). Averaging theory. In C. M. Bradshaw, E. Szabadi, & C. F. Lowe (Eds.), *Quantification of steady state operant behaviour* (pp. 21-24). Amsterdam: Elsevier.
- Logan, F. A., & Ferraro, D. P. (1978). *Systematic analyses of learning and motivation*. New York: Wiley.
- Marr, M. J. (1985). 'Tis the gift to be simple: A retrospective appreciation of Mach's *The Science of Mechanics*. *Journal of the Experimental Analysis of Behavior*, **44**, 129-138.
- Medawar, P. B. (1967). *The art of the soluble*. London: Methuen.
- Minsky, M. L. (1967). *Computation: Finite and infinite machines*. Englewood Cliffs, NJ: Prentice-Hall.
- Shimp, C. P. (1966). Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, **9**, 443-455.
- Skinner, B. F. (1935). The generic nature of the concepts of stimulus and response. *Journal of General Psychology*, **12**, 40-65.
- Skinner, B. F. (1969). *Contingencies of reinforcement: A theoretical analysis*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1972). A lecture on "having" a poem. In *Cumulative record: A selection of papers* (3rd ed., pp. 345-355). New York: Appleton-Century-Crofts.
- Smith, L. D. (1986). *Behaviorism and logical positivism: A reassessment of the alliance*. Stanford, CA: Stanford University Press.
- Staddon, J. E. R. (1967). Asymptotic behavior: The concept of the operant. *Psychological Review*, **74**, 377-391.
- Staddon, J. E. R. (1973). On the notion of cause, with applications to behaviorism. *Behaviorism*, **1**, 25-63.
- Staddon, J. E. R. (1984). Social learning theory and the dynamics of interaction. *Psychological Review*, **91**, 502-507.
- Staddon, J. E. R., & Bueno, J. L. O. (1991). On models, behaviorism and the neural basis of learning. *Psychological Science*, **2**, 3-11.
- Stevens, S. S. (1951). *Handbook of experimental psychology*. New York: Wiley.
- Verplanck, W. (1957). A glossary of some terms used in the objective science of behavior. *Psychological Review*, **64** (No. 6, Part 2).
- Williams, B. A. (1976). Short-term retention of response outcome as a determinant of serial reversal learning. *Learning and Motivation*, **7**, 418-430.